

Technical Report
828

AD-A230 061

Initialization for Improved IIR Filter Performance

E.S. Chornoboy

DTIC
ELECTE
JAN 23 1991
S E D

31 December 1990

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS



Prepared for the Federal Aviation Administration
under Air Force Contract F19628-90-C-0002.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

This document is disseminated under the sponsorship of the Federal Aviation Administration, Department of Transportation in the interest of information exchange. The U.S. Government assumes no liability for its contents or use thereof.

1. Report No. TR-828	2. Government Accession No. DOT/FAA/NR-90/7	3. Recipient's Catalog No.	
4. Title and Subtitle Initialization for Improved IIR Filter Performance		5. Report Date 31 December 1990	
		6. Performing Organization Code	
7. Author(s) Edward S. Chornoboy		8. Performing Organization Report No. TR-828	
9. Performing Organization Name and Address Lincoln Laboratory, MIT P.O. Box 73 Lexington, MA 02173-9108		10. Work Unit No. (TRAIS)	
		11. Contract or Grant No. DTFA-01-L-83-4-10579	
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration Systems Research and Development Service Washington, DC 20591		13. Type of Report and Period Covered Technical Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology under Air Force Contract F19628-90-C-0002.			
16. Abstract A new method for initializing the memory registers of Infinite Impulse Response (IIR) filters is presented. This method is shown to significantly reduce the initial transients which accompany the filtering of finite-length data sequences. Unlike previous methods, the proposed method makes no a priori assumptions regarding the input signal. Therefore, the method applies equally well to a variety of IIR designs and applications. The method does require a leading segment of the input data for initialization computations before filtering can begin. For this reason, the method is best suited for signal-processing applications in which batch processing of the data is employed. In particular, the method could prove very useful in situations where data is at a premium and only short-length sequences are available, because almost all data is usable after filtering. Applications using sequential processing of data can be accommodated when delays at the beginning of a processing segment can be tolerated.			
17. Key Words		18. Distribution Statement Document is available to the public through the National Technical Information Service, Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 19	22. Price

ABSTRACT

A new method for initializing the memory registers of Infinite Impulse Response (IIR) filters is presented. This method is shown to significantly reduce the initial transients which accompany the filtering of finite-length data sequences. Unlike previous methods, the proposed method makes no *a priori* assumptions regarding the input signal. Therefore, the method applies equally well to a variety of IIR designs and applications.

The method does require a leading segment of the input data for initialization computations before filtering can begin. For this reason, the method is best suited for signal-processing applications in which batch processing of the data is employed. In particular, the method could prove very useful in situations where data is at a premium and only short-length sequences are available, because almost all data is usable after filtering. Applications using sequential-processing of data can be accommodated when delays at the beginning of a processing segment can be tolerated.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input checked="checked" type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



TABLE OF CONTENTS

Abstract	ii
List of Illustrations	iv
1. INTRODUCTION	1
2. IIR TRANSIENTS AND INITIALIZATION	3
2.1 State-Variable Description	3
2.2 Complex-Exponential Input	4
2.3 Vector-Space Setting	8
2.4 Implementation	10
3. RESULTS	11
3.1 Initialization from Partial Observations	11
3.2 Average-Output-Magnitude Response	13
4. CONCLUSIONS	14
REFERENCES	15

LIST OF ILLUSTRATIONS

Figure No		Page
1	The average-output-magnitude frequency response of a four-pole elliptic IIR filter (after Fletcher and Burlage).	2
2	Conventional direct-form implementation of a second-order IIR digital filter.	3
3	Transfer-function magnitudes for a second-order (high-pass) IIR filter:	5
4	Filtering a short-length data sequence—zero-initialized transient-response characterization.	6
5	Filtering a short-length data sequence—step-initialized transient-response characterization.	7
6	Filtering a short-length data sequence—projection-initialized transient-response characterization.	11
7	Initialization using partial observations—spectral magnitude at system resonance vs. input frequency.	12
8	The average-output-magnitude frequency response of a four-pole elliptic IIR filter—projection-initialized results.	13

1. INTRODUCTION

Comparisons between Finite Impulse Response (FIR) and Infinite Impulse Response (IIR) filters are often based on the computational complexity¹ required to achieve a given frequency response. For applications requiring filters with sharp cutoffs (narrow transition regions), IIR filters are generally considered better because of the large filter length required in an FIR implementation. This distinction also occurs with (or is compounded by) applications in which the unfiltered data sequence has a short overall length or comprises short-length non-contiguous segments.² In such cases, FIR filters with sharp cutoffs can become impractical, requiring too great a proportion of the data for initialization of the filter taps. Because the all-zero portion of a comparable IIR filter is considerably shorter in length, an IIR filter could conceivably be used in situations with small-sized data sets. Unfortunately, the transient response of an IIR filter distorts the output of the filter on startup, rendering the situation similar to that of the FIR filter (i.e., an initial segment of data is consumed to initialize the filter; in the case of IIR filtering, output sampling is delayed until transients are significantly reduced). For an FIR filter, the amount of data lost is equal to the filter length N ; by one optimistic rule of thumb, a comparable IIR filter would require approximately the same N samples to significantly reduce its transients. To deal with this problem, it is often suggested that state initialization (i.e., initialization of the IIR memory registers) be used to reduce the IIR startup transients. The most common method, which is further described below, is based on approximating the unfiltered signal by a step input. This report presents a new alternative initialization method derived from a state-variable description and vector-space view of the transient problem. In contrast to the step-input approximation, this new method is equally suited to all type IIR structures (i.e., low-pass, bandpass, or high-pass).

Fletcher and Burlage, [2] and [3], working in the context of radar-clutter filtering³ for Moving Target Indicator (MTI) radar, proposed a method for improving the performance (i.e., reducing the transients) of IIR implementations in sampling situations as described above. However, their method is specific for IIR filters designed to filter low-frequency ground-clutter targets; it assumes that clutter-return signals can be modeled as step inputs with magnitude equal to that of the first recorded data sample (referred to as step initialization). Assuming that the ground-clutter step dominates the desired input signal, an appeal to the final-value theorem for Z-transforms can be used to calculate an approximation to the steady-state values of the filter memory registers. In the resulting procedure, filter memory registers are loaded with a scaled version of the leading data sample (the scaling factor is predetermined by the filter coefficients) prior to the normal filtering of the data. The method's simplicity is clearly an advantage, but Figure 1 illustrates the limited extent to which the filter's average-output-magnitude frequency response can be improved. As can be seen from the figure, the response obtained using the initialization method is still far from the desired steady-state response; in particular, the inability to match the extent of the stop-band width is a significant deficiency. Note that neither delay sampling nor delay sampling after initialization offered a substantial improvement toward approximating steady-state performance.

Development of the initialization method presented in this report was motivated by Lincoln Laboratory's work with high-pass filters and clutter removal for pulsed-Doppler radars. An initial objective was

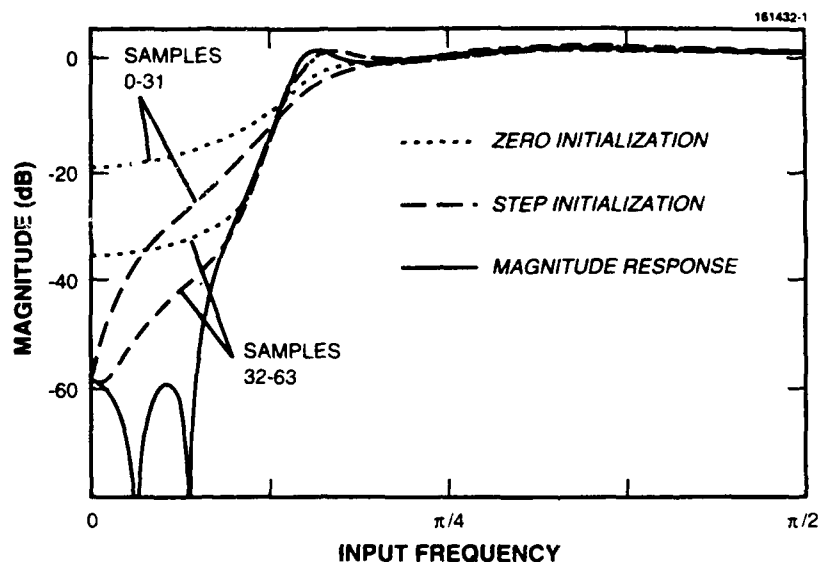


Figure 1. The average-output-magnitude frequency response of a four-pole elliptic IIR filter (after Fletcher and Burlage). The dotted lines illustrate average-magnitude responses for a zero-initialized filter; responses were computed by summing the magnitudes of output samples 0-31 and 32-63. Therefore, the lower of the two plots (from samples 32-63) represents the average response profile obtained after delay-sampling 32 samples. Dashed lines represent the corresponding profiles for the step-initialization filter. The solid curve shows the steady-state magnitude response of the filter. Responses are plotted for positive frequencies out to a frequency equal to one half the Nyquist value.

to improve upon the quantitative results obtainable by step initialization (Figure 1). However, this new method is considered to be of general interest because it makes no *a priori* assumptions regarding the content of the input signal or filtering application; the results apply equally well with any design IIR filter. However, the improvements realized with this new method are obtained at the cost of additional computations; in contrast to using only the first data point for initialization, the proposed method requires a leading segment of the input sequence for initialization computations.⁴ Hence, real-time applications with sequential processing must be amenable to delays at the start of a data block to acquire data for initialization. For real-time applications using batch processing, this is not a problem. The required computations are straightforward—dot products with real-valued weighting vectors—and the additional computations are not viewed as a prohibitive handicap.

2. IIR TRANSIENTS AND INITIALIZATION

Primary attention will be focused on the response of a general second-order filter, understanding that higher-order filters can be obtained by cascading. Extension of the solution to cascaded structures is outlined briefly at the section's end. All filter coefficients are assumed to be real valued; a block diagram for the conventional direct-form implementation, which is assumed, is provided for reference in Figure 2.

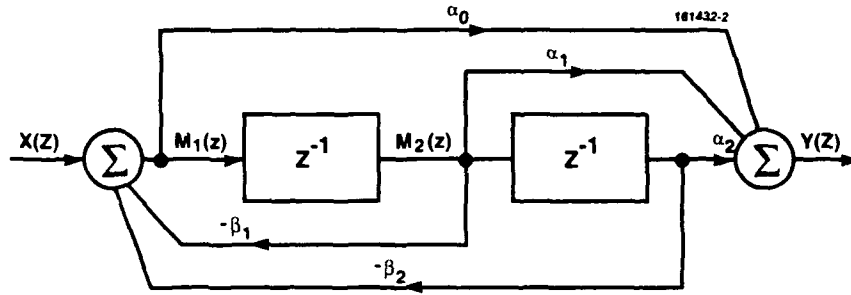


Figure 2. Conventional direct-form implementation of a second-order IIR digital filter: selecting starting values for the filter memory registers $m_1(n) \{ \tilde{z} M_1(z) \}$ and $m_2(n) \{ \tilde{z} M_2(z) \}$, the object of IIR initialization.

2.1 STATE-VARIABLE DESCRIPTION

For a general second-order filter, complex-valued input $x = \{x(n); n > -\infty\}$ and output $y = \{y(n); n > -\infty\}$ are related by the defining equation

$$y(n) = \alpha_0 x(n) + \alpha_1 x(n-1) + \alpha_2 x(n-2) - \beta_1 y(n-1) - \beta_2 y(n-2)$$

which correlates with the Z-transfer function definition

$$H(z) = \frac{\Delta Y(z)}{X(z)} = \frac{\alpha_0 + \alpha_1 z^{-1} + \alpha_2 z^{-2}}{1 + \beta_1 z^{-1} + \beta_2 z^{-2}}$$

The internal state of the filter section is specified by the vector

$$m(n) = \begin{bmatrix} m_1(n) \\ m_2(n) \end{bmatrix} \quad \left\{ \begin{array}{l} z \\ \leftrightarrow M(z) = \begin{bmatrix} M_1(z) \\ M_2(z) \end{bmatrix} \end{array} \right\}$$

(refer to Figure 2), and the following state-variable description of the filtering equations follows directly:

$$M(n) = B M(n-1) + C x(n) \quad (1a)$$

and

$$y(n) = A^T M(n-1) + \alpha_0 x(n) \quad (1b)$$

where

$$A \triangleq \begin{bmatrix} \alpha_1 - \alpha_0 \beta_1 \\ \alpha_2 - \alpha_0 \beta_2 \end{bmatrix}, \quad B \triangleq \begin{bmatrix} -\beta_1 - \beta_2 \\ 1 \quad 0 \end{bmatrix}, \quad \text{and } C \triangleq \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

With only a finite-windowed version of the data $\{x(n); 0 \leq n < N\}$ available, the state-transition equation is written

$$M(n) = B^n M_{-1} + \sum_{k=1}^n B^{n-k} C x(k), \quad (2)$$

where $M_{-1} = [m_1(-1) \ m_2(-1)]^T$ is the parameter describing the initial state of the filter section. As M_{-1} is generally undefined in applications, the goal of IIR initialization is selection of a value \hat{M}_{-1} that is appropriate for the filtering application.

The simplest form of initialization is to take $\hat{M}_{-1} = \emptyset$ (the null vector), which will be referred to as zero initialization. In step initialization, $\hat{M}_{-1} = \{1 + \beta_1 + \beta_2\}^{-1} [x(0) \ x(0)]^T$ is chosen to eliminate transient response due to zero-frequency energy.

2.2 COMPLEX-EXPONENTIAL INPUT

Working the problem for complex-exponential input allows solving directly for the transient contribution to the output response. A one-sided Z-transform analysis for a step-modulated¹ complex-exponential input, $x(n) = e^{j\omega n} u(n)$, yields the transient output response (assuming for now that $M_{-1} = \emptyset$)

$$y_{tr}(n) = -H_D(e^{j\omega}) A^T \left[I \delta(n) + r^{n-1} \frac{\sin n\theta}{\sin \theta} B u(n-1) - r^{n-2} \frac{\sin(n-1)\theta}{\sin \theta} \beta_2 I u(n-2) \right] \mathcal{E}(e^{j\omega}), \quad (3)$$

where

$$H_D(e^{j\omega}) = \frac{1}{(1 - re^{j(\theta-\omega)})(1 - re^{-j(\theta+\omega)})}, \quad \mathcal{E}(e^{j\omega}) = \begin{bmatrix} e^{-j\omega} \\ e^{-j2\omega} \end{bmatrix},$$

and I is the identity matrix. Note that in Equation (3) θ and r follow from a polar representation for the filter poles (i.e., $P = re^{\pm j\theta}$) and are related to the filter coefficients by $\beta_1 = -2r \cos \theta$ and $\beta_2 = r^2$ ([1], pp. 150-162). The vector $\mathcal{E}(e^{j\omega})$ can be viewed as an extension of input samples to times -1 and -2 . The gain $H_D(e^{j\omega})$ is from the magnitude response of the all-pole (or recursive) component of the filter. Because $H_D(e^{j\omega})$ is determined solely by the filter poles, signal transients can be disproportionately larger for input energy in the stop-band regions (see Figure 3).

Assuming the processing of short-length data sequences, it is helpful to have an example/characterization of the transient-response effect, both in the frequency and time domains, to help visualize the above complex exponential case as well as to provide a basis for later comparisons. This effect is

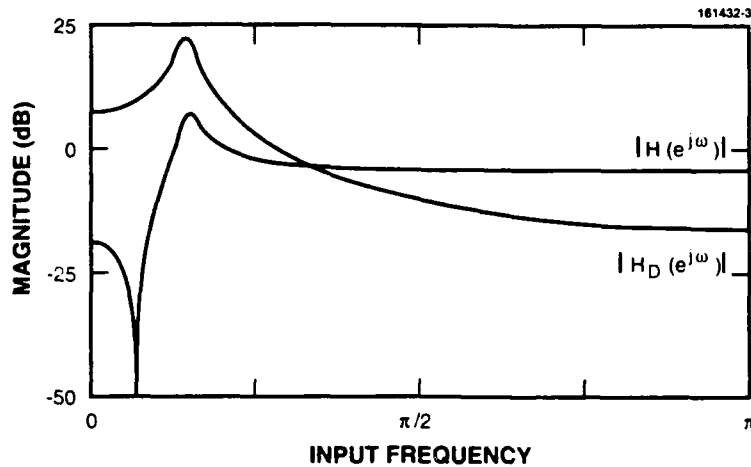


Figure 3. Transfer-function magnitudes for a second-order (high-pass) IIR filter: the frequency magnitude response, $|H(e^{j\omega})|$, and the denominator response, $|H_D(e^{j\omega})|$, are both plotted. The denominator response is the response due to the all-pole (recursive) portion of the filter.

introduced in Figures 4 and 5, which were derived from the processing of 64-sample complex-exponential sequences. Input signals were generated with frequencies spanning the range $0 - \pi$ radians. Each sequence was processed by a second-order high-pass IIR filter, whose magnitude response is plotted in Figure 3. The responses in Figure 4 result from zero-initialized processing of the sequences; those of Figure 5 from step-initialized processing. Each figure contains time-domain plots of the transient-component magnitude (left side) and periodogram spectral estimates of the entire output sequence (right side).

In the time-domain plots, oscillation near the resonant frequency and a characteristic r^n decay are both evident as functions of sample number. In the zero-initialized case, the magnitude of the transient response vs. input frequency (the axis going into the page) indicates modulation by $|H_D(e^{j\omega})|$. The step-initialized case has a similar characteristic, but shows a marked improvement at zero frequency and a significant magnitude increase for frequencies in the filter's passband.

The periodograms both indicate two primary spectral contributions, one a positive frequency (sweeping the range $0 - \pi$ and corresponding to the input-signal frequency) and the other a real sinusoidal component [energy at both $\pm \theta$, see Equation (3)], attributable to the filter's transient response. The characterization in this figure can be correlated with that presented in Figure 1, in which the average-output-magnitude frequency response measures the total output energy corresponding to a given sinusoidal input. Each spectral analysis includes a projection (onto the Magnitude vs. Input-Frequency plane) of a component attributable to the transient response. In the case of zero initialization, the shape of this component is due to weighting both by $|H_D(e^{j\omega})|$ and the geometric decay of r^n ; however, the influence of $|H_D(e^{j\omega})|$ is still apparent. For step initialization, this spectral component clearly illustrates improved performance at zero input frequency but illustrates an apparent compromise at higher passband frequencies.

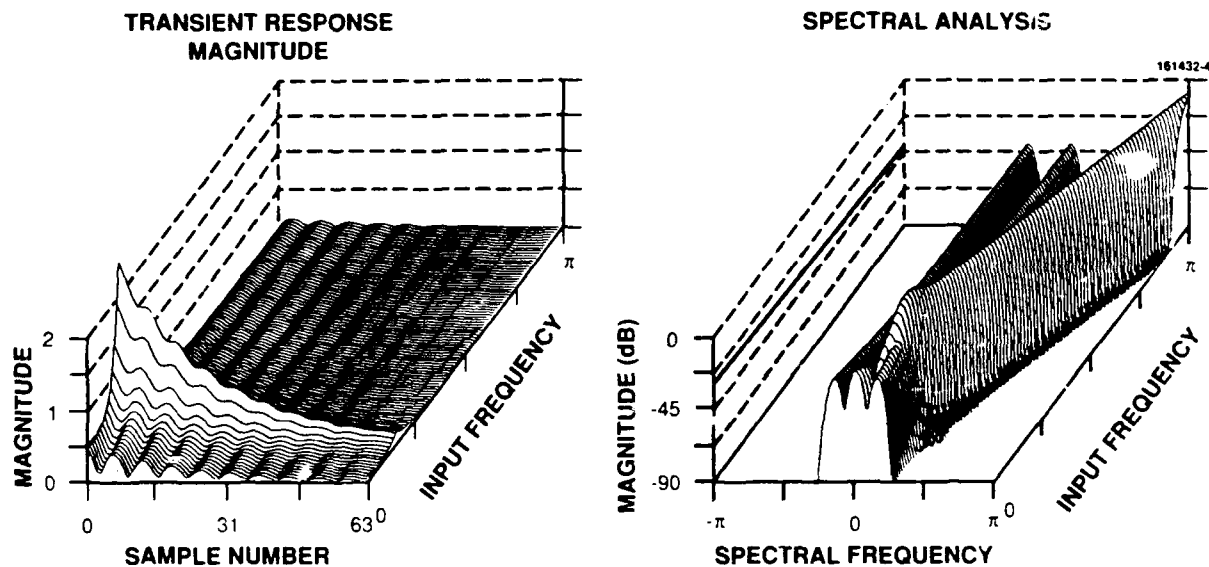


Figure 4. Filtering a short-length data sequence—zero-initialized transient-response characterization: time-domain values of the transient-component magnitude (left side) and windowed-periodogram (Kaiser weighting) spectral estimates of the complete 64-point output signal (right side) are plotted. Input signals were sampled complex exponentials having frequencies in the range 0 to π (plots are indexed with respect to input-signal frequency from foreground to background). Each input sequence was processed by a second-order IIR (high-pass) filter stage (see Figure 3). All filter memory registers were zeroed prior to filtering. Values for the transient-response magnitude are represented on a linear-magnitude scale; those for the spectral analysis, a log-magnitude scale (values are expressed in dB relative to an arbitrarily chosen level). The transient response introduces a signal component with energy measurable near system resonance. A shaded plot of the spectral component near resonance is projected onto the Magnitude vs. Input Frequency plane.

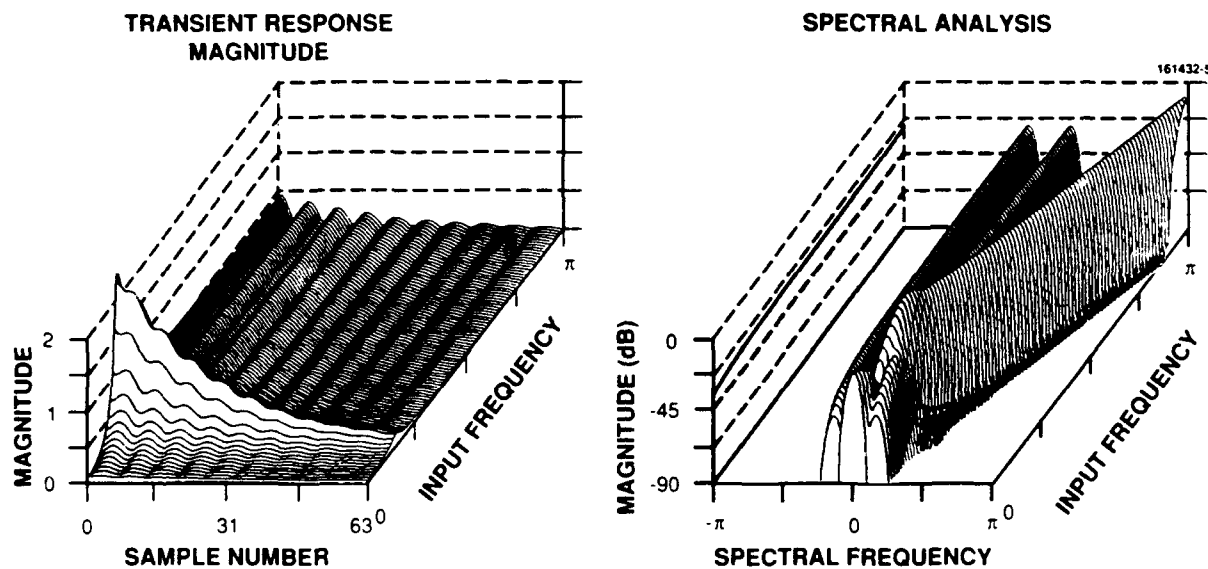


Figure 5. Filtering a short-length data sequence—step-initialized transient-response characterization: the plots in this figure are companion plots to those of Figure 4. The description of this figure is the same as in the previous figure, except that these plots result from processing the data with a step-initialized filter. Note: in contrast to the previous figure, there is no transient response corresponding to a zero-frequency input. However, this desired attenuation of the transient response is quickly lost as the input frequency increases. For input frequencies in the filter's passband, the magnitude of the transient response is actually greater than in the zero-initialized case.

Continuing the one-sided Z-transform analysis to compute the response to M_{-1} (i.e., $X \equiv \emptyset$, which also provides the homogeneous-equation solution), the same form as Equation (3) is obtained, as is the conclusion that $M_{-1} = H_D(e^{j\omega}) \mathcal{E}(e^{j\omega})$ yields $y_{tr}(n) \equiv 0$ for the case of complex-exponential input. As a special case, taking $\omega = 0$ results in $M_{-1} = \{1 + \beta_1 + \beta_2\}^{-1} [1 \ 1]^T$, which is the step-initialized solution.

The transient response to any step-modulated signal must have the form of Equation (3), because by direct analogy with the solution of second-order differential systems, the transient response must have the form of the homogeneous solution to the second-order difference equation. By viewing the transient component $\{y_{tr}(n); 0 \leq n < N\}$ as an element in an N -dimensional vector space, the conclusion is that the transient response (for a second-order system) is restricted to a two-dimensional subspace. This also can be seen in Equation (3) if $\mathcal{E}(e^{j\omega})$ is viewed as an arbitrary 2×1 weighting vector (representing the only degrees of freedom). Hence, for N sufficiently larger than 2, this constraint on the transient component may be used to advantage for initialization. This is the basis for the proposed initialization method, referred to as projection initialization.

2.3 VECTOR-SPACE SETTING

Let input and output sequences be viewed as $N \times 1$ vectors in a complex-valued vector space: $X, Y \in C$, with $Y = [y(0) \ y(1) \ \cdots \ y(N-1)]^T$ and $X = [x(0) \ x(1) \ \cdots \ x(N-1)]^T$. By componentwise application of Equations (1) and (2), linear operators $F (C^2 \rightarrow C^N)$ and $G (C^N \rightarrow C^N)$ can be defined which describe the relationship between input, output, and initial conditions:

$$Y = F M_{-1} + GX \quad (4)$$

Specifically, F is the $N \times 2$ matrix

$$F = \begin{bmatrix} A^T \\ A^T B \\ \dots \\ A^T B^{N-1} \end{bmatrix}$$

and G is the $N \times N$ lower triangular matrix

$$G = \begin{bmatrix} \alpha_0 & 0 & 0 & 0 & 0 & 0 \\ A^T C & \alpha_0 & 0 & 0 & 0 & 0 \\ A^T BC & A^T C & 0 & 0 & 0 & 0 \\ \cdot & \cdot & \alpha_0 & 0 & 0 & 0 \\ \cdot & \cdot & A^T C & \alpha_0 & 0 & 0 \\ \cdot & \cdot & A^T BC & A^T C & \alpha_0 & 0 \\ A^T B^{N-2} C & A^T B^{N-3} C & A^T B^2 C & A^T BC & A^T C & \alpha \end{bmatrix}$$

In the context of the initialization problem, the filter output can be viewed as being decomposed into steady-state and transient components: $Y = Y_{ss} + Y_{tr}$, where following from the previous section $Y_{tr} \in R(F)$ (the range of F) can be declared. Ideally, the goal of initialization would be to obtain $Y_{tr} = \emptyset$; or, equivalently, to find an appropriate criterion for minimizing $\|Y_{tr}\|$. Neither problem can be formulated directly; it is necessary to work an approximate problem. In particular, the objective is to make the transient distortion small, or at least reasonable for a practicable filter.

The projection operator [4] $P_F = F(F^T F)^{-1} F^T$ can be used to provide a decomposition for C^N , and the output response can be written as the sum of two orthogonal components: $Y = (I - P_F)Y + P_F Y$. The orthogonality of the decomposition provides the relation

$$\|Y\|^2 = \|(I - P_F)Y\|^2 + \|P_F Y\|^2$$

which, because P_F is a projector for the transient-response subspace, can also be written

$$\|Y\|^2 = \|(I - P_F)Y_{ss}\|^2 + \|P_F Y_{ss} + Y_{tr}\|^2$$

The first term on the right, $\|(I - P_F)Y_{ss}\|$, does not depend on M_{-1} . If the treatment of $\|P_F Y_{ss} + Y_{tr}\|$ could be rationalized as an approximation to $\|Y_{tr}\|$, an initializing \hat{M}_{-1} could be selected as the solution to

$$\hat{M}_{-1} = \arg \min_{M_{-1} \in C^2} \|P_F Y\| \quad (5)$$

Proposition 1. For the second-order IIR filtering problem described above, selecting

$$\hat{M}_{-1} = -(F^T F)^{-1} F^T G X \quad (6)$$

ensures that

$$\max_{0 \leq n < N} |y_{tr}(n)| \leq \|Y_{tr}\| \leq \left\{ \frac{1}{2\pi} \int_{-\pi}^{\pi} |H(e^{j\omega})|^2 |X(e^{j\omega})|^2 d\omega \right\}^{\frac{1}{2}}$$

where $\|\cdot\|$ represents the normal Euclidean norm.

In the case of complex-exponential input, the transient-component magnitude will be proportional to the filter's overall magnitude response $|H(e^{j\omega})|$ instead of the denominator magnitude response $|H_D(e^{j\omega})|$. The most significant consequence is that transients elicited by input-signal energy in the stop band are greatly attenuated, because the filter's zeros now play a role in determining the transient response. The next section continues the previous numerical examples and illustrates the improved stop-band performance.

Proof. The solution to Equation (5) is obtained by solving $P_F Y = 0$, which, after substituting Equation (4) and rearranging, becomes $F M_{-1} = -P_F G X$. Equation (6) follows, given the definition of P_F . For this choice of M_{-1} we have $Y_{tr} = -P_F Y_{ss}$; therefore, using Parseval's relation,

$$\|y_{tr}\| \leq \|Y_{ss}\| = \left\{ \frac{1}{2\pi} \int_{-\pi}^{\pi} |H(e^{j\omega})|^2 |X(e^{j\omega})|^2 d\omega \right\}^{\frac{1}{2}}$$

(P_F being a projection operator and therefore idempotent, has a spectral norm of 1).

2.4 IMPLEMENTATION

Implementation of Equation (6) is straightforward, as the matrix $(F^T F)^{-1} F^T G$ is completely predetermined by knowledge of the filter coefficients. The coefficients of the weighting matrix are all real valued; the initialization computations only require real-valued dot products between input data and vectors of weighting coefficients. For a second-order filter, four such multiplications are required; one each for the real and imaginary components of both memory registers. Once an initialization vector has been computed, the initialized IIR filter can be implemented using standard methods. Alternatively, the output Y can also be computed directly by implementing the matrix multiplication $Y = (I - P_F)GX$.

Higher-Order Filters: Cascading. For a cascade of second-order sections, repeated application of Equation (4) can derive a relationship between input and output that is analogous to Equation (4): $Y = \tilde{F}\tilde{M}_{-1} + \tilde{G}X$, where

$$\begin{aligned}\tilde{F} &\triangleq \left[\left(\prod_{i=2}^n G_i \right) F_1 \mid \left(\prod_{i=3}^n G_i \right) F_2 \mid \cdots \mid G_n F_{n-1} \mid F_n \right], \\ \tilde{M}_{-1} &\triangleq \left[M_{-1,1}^T \mid M_{-1,2}^T \mid \cdots \mid M_{-1,n}^T \right]^T, \\ \tilde{G} &\triangleq \prod_{i=1}^n G_i, \\ \prod_{i=k_1}^{k_2} G_i &\triangleq G_{k_2} G_{k_2-1} \cdots G_{k_1},\end{aligned}$$

and F_i , G_i , and $M_{-1,i}$ represent the coefficient matrices and initialization vector for the i^{th} filter stage. Note that now \tilde{F} is $N \times 2\eta$ and \tilde{M}_{-1} is $2\eta \times 1$, where η is the number of second-order stages in the filter. Correspondingly, the dimension of $R(\tilde{F})$ will be 2η (assuming no degenerate cases). The initialization solution here is wholly analogous to the second-order case, and the solution [Equation (6)] applies with the substitution of \tilde{M}_{-1} , \tilde{F} , and \tilde{G} for M_{-1} , F , and G .

3. RESULTS

The previous numerical example is completed by including the case of a projection-initialized filter, i.e., a filter initialized by Equation (6). The characterization is presented in Figure 6. Referring to the time-domain plots, with projection initialization there is considerable improvement in the stop-band extent of transient reduction; there is also a reduction in transient magnitude for much of the passband region as well. Only for input-signal energy near the resonant frequency is there an obvious increase in transient magnitude. Depending upon the application, this increase may not be significant, as the transient component is near the resonant frequency, giving rise to a spectral component corresponding to the input frequency. The accompanying periodogram plots clearly illustrate the motivation for Equation (6). Viewing the magnitude of the transient-response spectral contribution (vs. input frequency), a modulation by $|H(e^{j\omega})|$ is clearly evident.

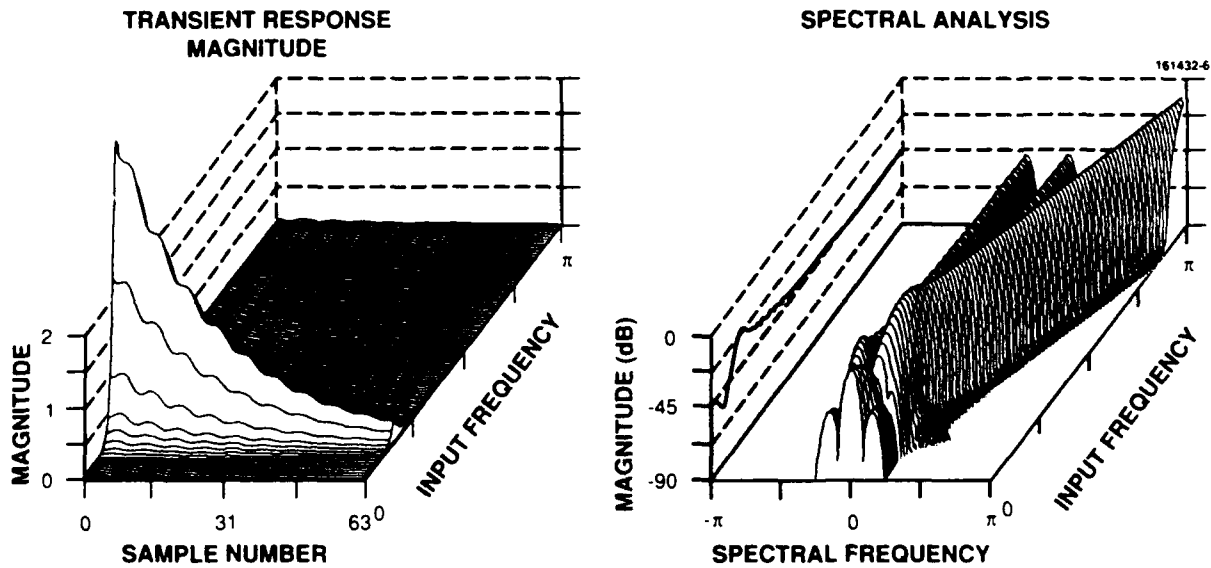


Figure 6. Filtering a short-length data sequence—projection-initialized transient-response characterization: the description of this figure is the same as that of Figure 4, except that the data were processed with a projection-initialized filter. In comparison to Figure 5, there is significant attenuation of the transient-component magnitude throughout most of the stop band and for a large portion of the passband as well. The magnitude of the transient-response spectral component indicates modulation by the filter's frequency-response magnitude. This is in contrast to the two previous cases where the denominator response modulates the transient-response energy.

3.1 INITIALIZATION FROM PARTIAL OBSERVATIONS

To this point it has been assumed that all data to be filtered are used for initialization; this need not be the case. For data processed sequentially (as opposed to batch) this usually cannot be achieved.

However, if the sequential processing can afford a delay of a number of data samples, initialization based on these data can be used. Figure 7 summarizes the effect of initialization sample size by plotting magnitudes of the transient-component spectral contributions for partial observation sets of 64, 48, 32, and 16 samples. The top plot for 64 samples, the entire data set, corresponds to the shaded (oblique) plot in Figure 6. Each panel also contains a plot of the corresponding magnitude for the zero-initialized case (see Figure 4). These results show that the method is still exceptionally effective at reducing the transients in response to input energy near the filter zero.

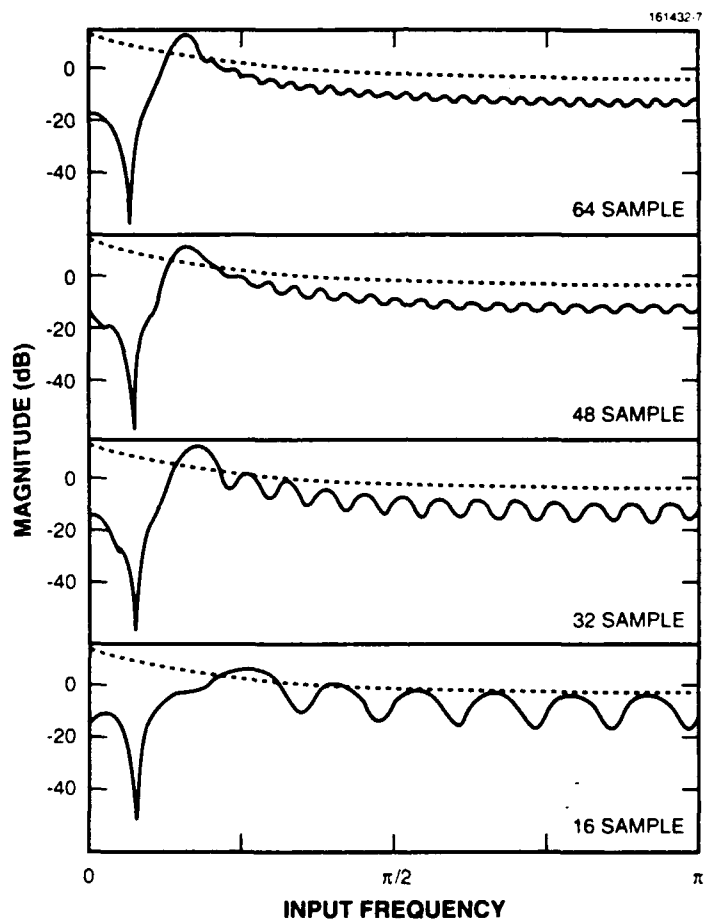


Figure 7. Initialization using partial observations—spectral magnitude at system resonance vs. input frequency: spectral magnitude at resonance is plotted vs. input frequency. The plots are derived, as in Figure 6, from projection-initialized high-pass filtering of 64-point input sequences. Four panels are illustrated corresponding to using the first 64, 48, 32, or 16 samples for initialization computations. The top panel, 64 samples, uses all available data and is the same as the oblique projection shown in Figure 6. For reference, each panel also contains a plot of the spectral magnitude corresponding to zero-initialized processing. Attenuation of the transient response for frequencies near the filter zero is well preserved.

3.2 AVERAGE-OUTPUT-MAGNITUDE RESPONSE

Referring to the four-pole example of Figure 1, a plot of the average-output-magnitude response for the projection-initialized filter is presented. Figure 8 reproduces the magnitude response of the filter and, for reference, the step-initialized results. Clearly, the projection-initialized filter demonstrates a much improved performance throughout the stop-band region. In contrast, however, there is a greater rounding of the transition to the passband, due to increased cancelling in the region of filter resonance.

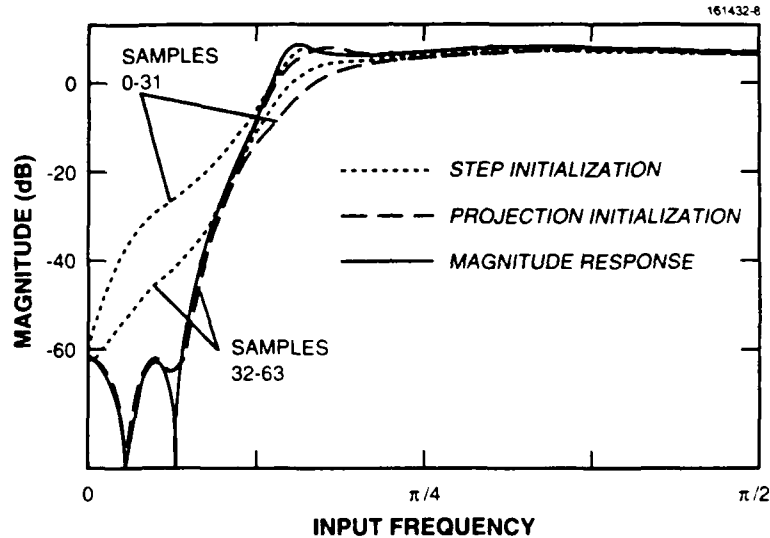


Figure 8. The average-output-magnitude frequency response of a four-pole elliptic IIR filter—projection-initialized results: the example of Figure 1 is completed with the results for projection-initialized filtering. The figure description is the same as in Figure 1, except that the zero-initialized results are not repeated and dotted lines are now used for the step-initialized results. The projection-initialized results, dashed lines, follow the steady-state stop-band performance. The response near filter resonance is more rounded, however, due to a larger transient response and cancellation at these frequencies (see Figure 6).

4. CONCLUSIONS

The transient response that occurs at the start of an IIR filtering pass usually results in the discarding of usable data to allow for IIR transient decay. In practical applications, much of the difficulty with transients stems from their disproportionately large magnitude in response to input energy in stop-band regions. This occurs because the transient response magnitude is determined by the filter poles. Projection initialization, which takes into account the location of the filter zeros, was shown to significantly reduce the transient-response magnitude in response to input energy near the filter zeros. Hence, regarding the filtering of short-length data sequences, the method can improve the effective stop-band performance of an IIR filter and substantially reduce the amount of data discarded.

The derivations and comparisons presented are based on a standard second-order filter definition; extension to higher-order filters was demonstrated by cascading. This does not imply that the method only works for even-ordered filter structures. Odd-ordered filters can be dealt with by placing the odd stage in second-order form (with some degenerate coefficients). The solutions presented can apply if a generalized inverse interpretation is applied to the solution.

Because the method requires an initial segment of data, it is ideally suited for applications where data is processed in batch mode. Sequential data processing can be accommodated only when delays at the beginning of a processing block can be tolerated. However, the method can improve performance, using only a small initial segment for the initialization computations.

REFERENCES

1. M. Bellanger, *Digital Processing of Signals: Theory and Practice*, John Wiley and Sons (1984).
2. R.H. Fletcher and D.W. Burlage, "An initialization technique for improved MTI performance in phased array radars," in *Proceedings of the IEEE* 60, 1551-1552 (1972).
3. R.H. Fletcher and D.W. Burlage, "Improved moving-target-indicator filtering for phased array radars," U.S. Army Missile Command, Redstone Arsenal, Alabama, Technical Report RE-73-17 (1973).
4. A.S. Householder, *The Theory of Matrices in Numerical Analysis*, New York: Dover Publications, Inc. (1975).